IRT-Lab

Version 2

User's Manual

CONTENTS

CHAPTER 1
Getting Started with IRT-Lab

IRT-Lab is a Windows based program that performs a variety of functions related to item response theory (IRT). These functions are aimed at facilitating research that use IRT-generated data, and improving the teaching and understanding of various IRT concepts.

System Requirements

The following hardware and software is required for running IRT-Lab Version 2:

- Microsoft Windows 95 or later, or Microsoft Windows NT 3.51 or later
- 24 MB RAM for Windows 95 or later, 32 MB RAM for Microsoft NT or later
- Pentium 90MHz or higher

Installation and Running Instructions

IRT-Lab consists of a single file that is approximately 500 kB in size. To run IRT-Lab, you must double-click on the IRT-Lab executable file, called "IRTLab2.exe". You may either run IRT-Lab directly from the disk on which it is distributed, or you may copy the executable file to your computer's hard drive. To copy the executable file, select the file "IRTLab.exe", select "Copy" from the "Edit" menu of Windows (or use the shortcut of Control-C), go to the directory where you want to paste the file and select "Paste" from the "Edit" menu of Windows (or use the shortcut of Control-V). The program will tend to run faster if it is run from your computer's hard drive than from the disk on which it is distributed.

Organization of IRT-Lab

IRT-Lab consists of three primary components: (1) the Main Window, (2) Instructional Windows, and (3) a Simulate Data Window. The Main Window includes components that permit the user to create a test of up to 100 items for up to 4 groups of respondents. The Instructional Windows consist of a set of Windows that illustrate various aspects of IRT parameter estimation, including the display of item characteristic curves (ICCs), likelihood functions related to item and person parameter estimation, and the display of item and test information functions. The Simulate Data Window permits the user to simulate data according to the parameters of the test items specified in the Main Window.

CHAPTER 2
Main Window

The Main Window contains components that permit you to create a test of up to 100 items for four different groups of individuals. This chapter describes how to construct a test, and assign the items of the test parameters.

Components of the Main Window

The left side of the Main Window consists of tools that permit the user to specify the properties of the test and items (e.g., the number of items in the test, and the parameters of the items). In all cases, the user specifies the parameters of the items according to one of six psychometric models:

> 1. Dichotomous Rasch model
> 2. One-parameter logistic model (1PL)
> 3. Two-Parameter logistic model (2PL)
> 4. Three-parameter logistic model (3PL)
> 5. Generalized partial credit model (GPCM)
> 6. Graded response model (GRM)

The equations associated with these models are presented in Chapter 5. Note that the test may be constructed using any combination of the above models. In addition, to mimic testing situations in which two or more distinct groups of examinees (e.g., different genders, races, ethnicities, countries, etc.) are administered the same test, and the items of the test function differentially across the two groups (that is, DIF exists for one or more items), the user may construct a test for which the item parameters differ between the groups. In total, IRT-Lab permits the construction of a test containing up to 100 items for up to 4 groups of examinees.

The right side of the Main Window consists of the Parameter Table, which displays the parameters of the items of the test. The headings of the Parameter Table are as follows:

```
Item   Grp   Mod    c     a      b1     b2     b3     b4     b5     b6
```

Note that there are a total of 11 columns to the Parameter Table. The columns of the test construction grid specify the following information:

Column 1:   Item number

The value in each row is specified according to the number of items specified in the associated drop-down box.

Column 2:   Group number

The value in each row is specified according to the number of groups specified in the associated drop-down box.

Column 3:   Model

The possible models include Rasch, 1PL, 2PL, 3PL, GPCM, and GRM.

Column 4   c-parameter

The user specifies the value of the c-parameter implicated in the 3PL. If the Rasch, 1PL, 2PL, 3PL, GPCM, or GRM are selected, the value in Column 4 is automatically set to 0.

Column 5   a-parameter

The user specifies the value of the a-parameter implicated in the 2PL, 3PL, GPCM, and GRM. If the Rasch or 1PL are selected, the value is automatically set to 1.

Column 6   b1-parameter

The user specifies the value of the b1-parameter. For the Rasch, 1PL, 2PL, 3PL models, the value corresponds to the difficulty parameter of the item. For the GPCM and GRM, b1 is the value of the location parameter for category 1, where the categories are denoted by 0, 1, …, m.

Column 7   b2-parameter

The user specifies the value of the b2-parameter for the GPCM and GRM having three or more categories.

Column 8   b3-parameter

The user specifies the value of the b3-parameter for the GPCM and GRM having four or more categories.

Column 9   b4-parameter

The user specifies the value of the b4-parameter for the GPCM and GRM having five or more categories.

Column 10   b5-parameter

The user specifies the value of the b5-parameter for the GPCM and GRM having six or more categories.

Column 11   b6-parameter

The user specifies the value of the b6-parameter for the GPCM and GRM having seven categories.

An example of the Parameter Table for a five-item test for a single group follows.

Item	Grp	Mod	c	a	b1	b2	b3	b4	b5	b6
1	1	Rasch	0	1	-1					
2	1	1PL	0	1	2.3					
3	1	2PL	0	0.5	0.2					
4	1	3PL	0.2	1.8	1.5					
5	1	GRM	0	0.9	-1.8	0.1	2.1			
6	1	GPCM	0	1.5	-2.1	-0.7	0.3	0.8	2.1	3.0

In this example, the first item follows the dichotomous Rasch model, the second item follows the 1PL, the third item follows the 2PL, the fourth item follows the 3PL, the fifth item follows the GRM with four response options, and the sixth item follows the GPCM with seven response options (note that for the GPCM and the GRM, an item with m response options is modeled using $m - 1$ location parameters).

Consider next an example for which there is a three-item test, where different parameters are permitted for two groups of examinees. An example of the Parameter Table for this situation is as follows.

Item	Grp	Mod	c	a	b1	b2	b3	b4	b5	b6
1	1	3PL	0.2	1.2	-1					
1	2	3PL	0.2	1.2	-0.8					
2	1	3PL	0.15	0.5	0.2					
2	2	3PL	0.15	1.2	0.2					
3	1	3PL	0.2	0.9	1.8					
3	2	3PL	0.2	1.5	2.1					

In this example, each item is duplicated for each group. Note, however, that the parameters for each item vary between the two groups. For Item 1, the difficulty parameter (b1) differs between the groups, for Item 2 the discrimination parameter (a) differs between the groups, and for Item 3 the difficulty and the discrimination parameters differ between the groups.

Creating a Test

The creation of a test begins by selecting the number of items and number of groups from the drop-down lists at the top left corner of the Main Window. By default, the number of groups is set equal to one, so you need not select the number of groups if you are concerned with the parameters for a single group. The maximum number of items that can be selected is 100, and the maximum number of groups that can be selected is four. Upon selecting items or groups, rows in the Parameters Table are appropriately created. For example, if three items and one group are selected, three rows will be created, one row for each item, and the Parameter Table will look like:

```
Item   Grp    Mod    c      a      b1     b2     b3     b4     b5     b6
1      1
2      1
3      1
```

Similarly, if three items and two groups are selected, six rows will be created, two rows for each item.

```
Item   Grp    Mod    c      a      b1     b2     b3     b4     b5     b6
1      1
1      2
2      1
2      2
3      1
3      2
```

Once the number of items and groups are selected, you may parameterize the items. To parameterize an item of the test, follow these steps:

Step 1: Select one or more items from the Parameter Table. To select a single item, left-click on the row of the Parameter Table corresponding to the item of interest. To select two or more adjacent items, left-click and drag the mouse pointer across the desired range of items.

Step 2: Select the desired model from the model drop-down list located below the item and group drop-down lists at the far left side of the Main Window.

Step 3: Specify the desired parameters of the selected items. Note that IRT-Lab will automatically correct any inappropriate parameter values. For example, if a value of the c-parameter of 0.2 is entered for a 1PL model, IRT-Lab will correct it to 0 when placing the parameters in the Parameter Table.

Step 4: Click on the "Update Parameters" command button. Upon doing so, the resulting parameters will appear in the Parameter Table.

If you have not provided IRT-Lab with appropriate information, a message will appear telling you which information is still required. In addition, it is not necessary to do Step 1 before Steps 2 and 3; you can do Steps 2 and 3 first, and then Step 1.

Using the Clear Command

The user may clear the contents of the Parameter Table using the Clear Command. To activate this command, select Clear from the Edit Menu. Activating this command will clear all contents of the Parameter Table, as well as resetting the number of groups to 1 and the number of items to zero.

Saving Item Parameters for Future Use

The user may save the contents of the Parameter Table to a file for future use in IRT-Lab. To save the contents of the Parameter Table, select "Save Parameters for Future IRT-Lab Session" from the File Menu of the Main Window. A dialogue box will appear that enables the user to set the desired name and location of the parameter file. The parameter file is saved with the extension ".irt".

Saving Item Parameters in a Text Log

The user may save the contents of the Parameter Table to a text file that may serve as a log of the parameters used in the current IRT-Lab session. To save the contents of the Parameter Table to a text log, select "Save Parameters in a Text Log" from the File Menu of the Main Window. A dialogue box will appear that enables the user to set the desired name and location of the text file. The text file is saved with the extension ".txt". This file many then be opened in any generic word processing program, such as Notepad or Word, and be printed from the word processing program. To maintain the proper formatting you must:

> 1. Change the font type to Courier or New Courier
>
> 2. Change the font size so that the lines do not wrap. Typically a font size of 9 is sufficient, although the appropriate size will likely vary depending on the program used and the settings of the program.

Opening Item Parameters

To open a file of item parameters created in IRT-lab, select Open Parameters from the File Menu of the Main Window. A dialogue box will appear that enables the user to select the desired IRT-Lab parameter file. Note that IRT-Lab parameter files have the extension ".irt".

Example 1: Creating a Test of Dichotomous Items

Step 1: Create a Test of Three Items. In the Item Drop Down Box at the top left corner of the IRT-Lab screen select "# Items = 3". The Item Parameters Table should now have three rows corresponding to three items. Notice that the Group Column reads "1" for each of the three items listed in the Item Parameters Table, indicating that there is only one group and that the item parameters correspond to those of the single group (as opposed to the situation when items function differentially for different groups, and thus different parameters are required for each group). The Parameter Table should now look like:

```
Item   Grp   Mod     c      a      b1     b2     b3     b4     b5     b6
1      1
2      1
3      1
```

Step 2: Set the model. In the model drop-down box, select 1PL.

Step 3: Set the parameter of Item 1. Left-click on the row of the Parameter Table corresponding to Item 1. The row corresponding to Item 1 should now be highlighted in blue. In the "b1" box, enter "-1". Click the "Update" button. The Parameter Table should now look like:

```
Item   Grp   Mod     c      a      b1     b2     b3     b4     b5     b6
1      1     1PL     0      1      -1
2      1
3      1
```

Step 4: Set the parameter of Item 2. Left-click on the row of the Parameter Table corresponding to Item 2. The row corresponding to Item 1 should now be highlighted in blue. In the "b1" box, enter "0". Click the "Update" button. The Parameter Table should now look like:

```
Item   Grp   Mod     c      a      b1     b2     b3     b4     b5     b6
1      1     1PL     0      1      -1
2      1     1PL     0      1      0
3      1
```

Step 5: Set the parameter of Item 3. Left-click on the row of the Parameter Table corresponding to Item 2. The row corresponding to Item 1 should now be highlighted in blue. In the "b1" box, enter "1". Click the "Update" button. The Parameter Table should now look like:

```
Item   Grp   Mod     c      a      b1     b2     b3     b4     b5     b6
1      1     1PL     0      1      -1
2      1     1PL     0      1      0
3      1     1PL     0      1      1
```

You have now created a test consisting of three items following the 1PL model. The difficulty parameters of the three items are –1, 0, and 1.

Example 2: Creating a Test of Polytomous Items

Step 1: Create a Test of Three Items. In the Item Drop Down Box at the top left corner of the IRT-Lab screen select "# Items = 3". The Item Parameters Table should now have three rows corresponding to three items. Notice that the Group Column reads "1" for each of the three items listed in the Item Parameters Table, indicating that there is only one group and that the item parameters correspond to those of the single group (as opposed to the situation when items function differentially for different groups, and thus different parameters are required for each group). The Parameter Table should now look like:

Item	Grp	Mod	c	a	b1	b2	b3	b4	b5	b6
1	1									
2	1									
3	1									

Step 2: Set the model. In the model drop-down box, select GPCM.

Step 3: Set the parameter of Item 1. Left-click on the row of the Parameter Table corresponding to Item 1. The row corresponding to Item 1 should now be highlighted in blue. In the "a" box enter "0.8", in the "b1" box enter "-2", in the "b2" box enter "0", and in the "b3" box enter "1". Click the "Update" button. The Parameter Table should now look like:

Item	Grp	Mod	c	a	b1	b2	b3	b4	b5	b6
1	1	GPCM	0	0.8	-2	0	1			
2	1									
3	1									

Step 4: Set the parameter of Item 2. Left-click on the row of the Parameter Table corresponding to Item 2. The row corresponding to Item 1 should now be highlighted in blue. In the "a" box enter "0.8", in the "b1" box enter "-1", in the "b2" box enter "1", and in the "b3" box enter "2". Click the "Update" button. The Parameter Table should now look like:

Item	Grp	Mod	c	a	b1	b2	b3	b4	b5	b6
1	1	GPCM	0	0.8	-2	0	1			
2	1	GPCM	0	0.8	-1	1	2			
3	1									

Step 5: Set the parameter of Item 3. Left-click on the row of the Parameter Table corresponding to Item 2. The row corresponding to Item 1 should now be highlighted in blue. In the "a" box enter "1.2", in the "b1" box enter "-1", in the "b2" box enter "1", and in the "b3" box enter "2". Click the "Update" button. The Parameter Table should now look like:

```
Item    Grp    Mod     c      a       b1      b2      b3      b4      b5      b6
1       1      GPCM    0      0.8     -2      0       1
2       1      GPCM    0      0.8     -1      1       2
3       1      GPCM    0      1.2     -1      1       2
```

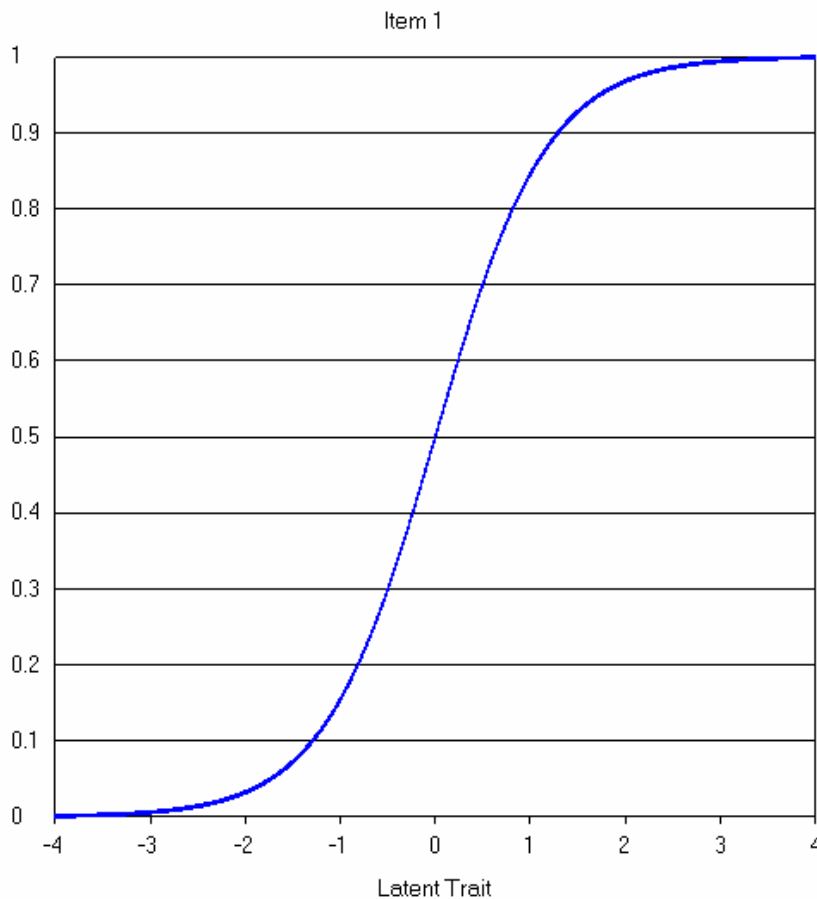You have now created a test consisting of three items following the GPCM.

CHAPTER 3
Instructional Windows

IRT-Lab includes several instructional windows that can be used to illustrate a variety of concepts related to IRT parameter estimation.
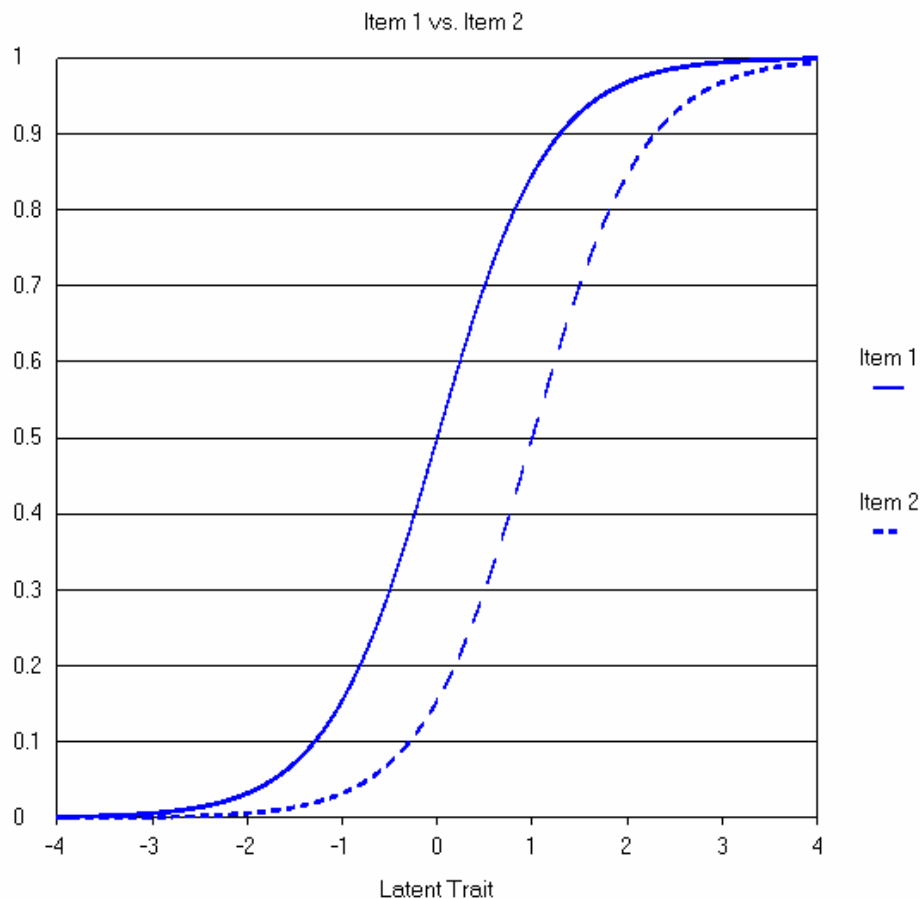
View ICCs

The View ICCs Window can be used to view the item characteristic curve (ICC) or curves (for polytomous items) for any of the items on the test. To view an item's ICC, select the desired item from the drop-down list at the top of the View ICCs Window. The resulting ICC will appear in the picture box of the View ICCs Window. An example of the resulting image for a 1PL item having $b1 = 0$ is as shown below.

Item 1

Latent Trait

The image can be copied and pasted into other applications by clicking on the "Copy" button at the right side of the View ICCs Window.

Compare ICCs

It is often desirable to view the ICCs of two item simultaneously. The Compare ICCs Window accomplishes this. To view the ICCs of two items simultaneously, select the first item from the "Select First Item of Comparison" drop-down list and the second item from the "Select Second Item of Comparison" drop-down list. The ICC of the first item is denoted by a solid line, and the ICC of the second item is denoted by a dashed line. An example of the resulting image for two items following the 1PL model (Item 1 has $b1 = 0$ and Item 2 has $b1 = 1$) is shown below.

The resulting image can be copied and pasted into other applications by clicking on the "Copy" button at the right side of the Compare ICCs Window.
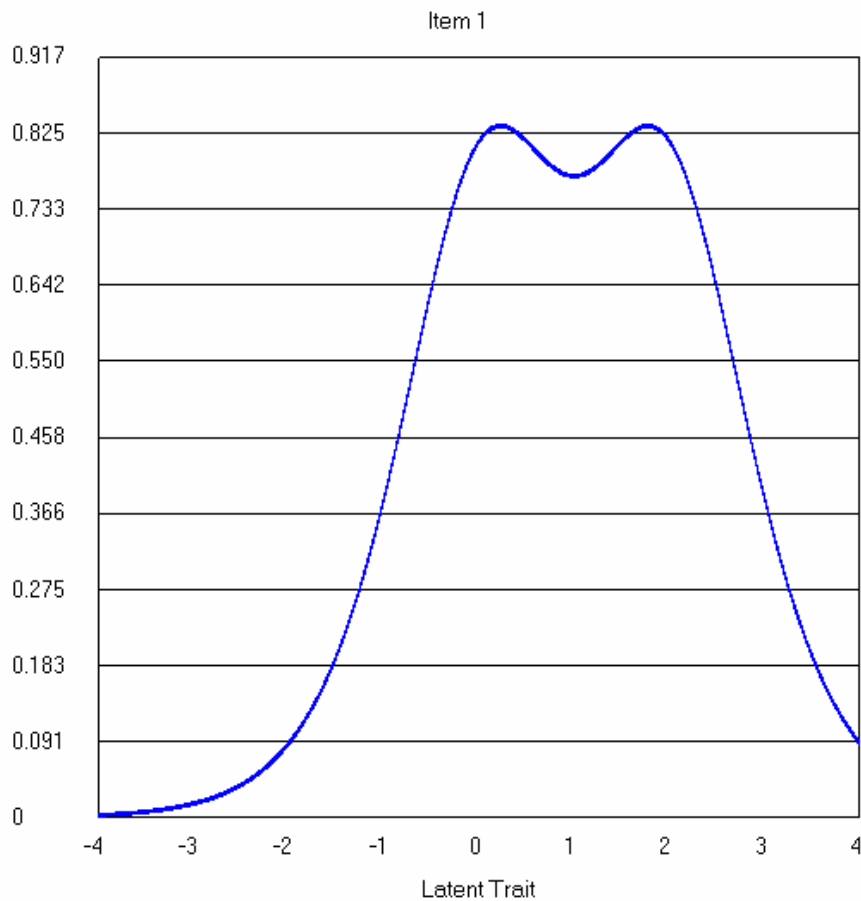
Item Estimation

The Item Estimation Window illustrates the maximum likelihood estimation of the difficulty parameter of a dichotomous item following a 1PL model. The user specifies the sample size (between 20 and 500), the actual value of the difficulty parameter (between –3 and 3) using the drop-down boxes at the top left corner of the window. When the user clicks the "Generate Responses" button, IRT-Lab simulates responses to the studied item for each individual in the selected sample, computes the log-likelihood function for the difficulty parameter, conducts a maximum likelihood estimation of the difficulty parameter, and displays the final estimate of the item difficulty in a box designated "Estimated b". The simulated responses are displayed in the Observed Scores Table. The Iterations and Estimates Table displays the estimated value of the difficulty parameter at each iteration of the Newton-Raphson algorithm, as well as the change in the value of the estimated item difficulty parameter at each iteration of the Newton-Raphson algorithm. The Newton-Raphson iterations are continued until the change in the difficulty parameter estimate is either less than the tolerance of 0.005 from the previous iteration, or a total of 25 iterations are completed. The log-likelihood function is displayed in a graph in the lower portion of the Item Estimation Window. The resulting image can be copied and pasted into other applications by clicking on the "Copy" button at the right side of the Item Estimation Window.

Theta Estimation

The Theta Estimation Window illustrates the maximum likelihood estimation of the latent trait level of an examinee responding to the items of the test specified in the Main Window. The user specifies the group number (between 1 and 4), the actual value of theta (θ) for the respondent in question (between –3 and 3) using the drop-down boxes at the top left corner of the Theta Estimation Window. When the user clicks the "Go" button, IRT-Lab simulates responses to the test items for the selected respondent, computes the log-likelihood function for θ, conducts a maximum likelihood estimation of θ, and displays the final estimate of θ in a box designated "Estimated Theta". The simulated responses are displayed in the Observed Scores Table. The log-likelihood function is displayed in the Graphic Box in the lower portion of the Theta Estimation Window. The Iterations and Estimates Table displays the estimated value of θ at each iteration of the Newton-Raphson algorithm, as well as the change in the value of the θ estimate at each iteration of the Newton-Raphson algorithm. The Newton-Raphson iterations are continued until the change in the estimate of θ is less than the tolerance of 0.005 from the previous iteration, or a maximum of 25 iterations of completed. The resulting image can be copied and pasted into other applications by clicking on the "Copy" button at the right side of the Theta Estimation Window.
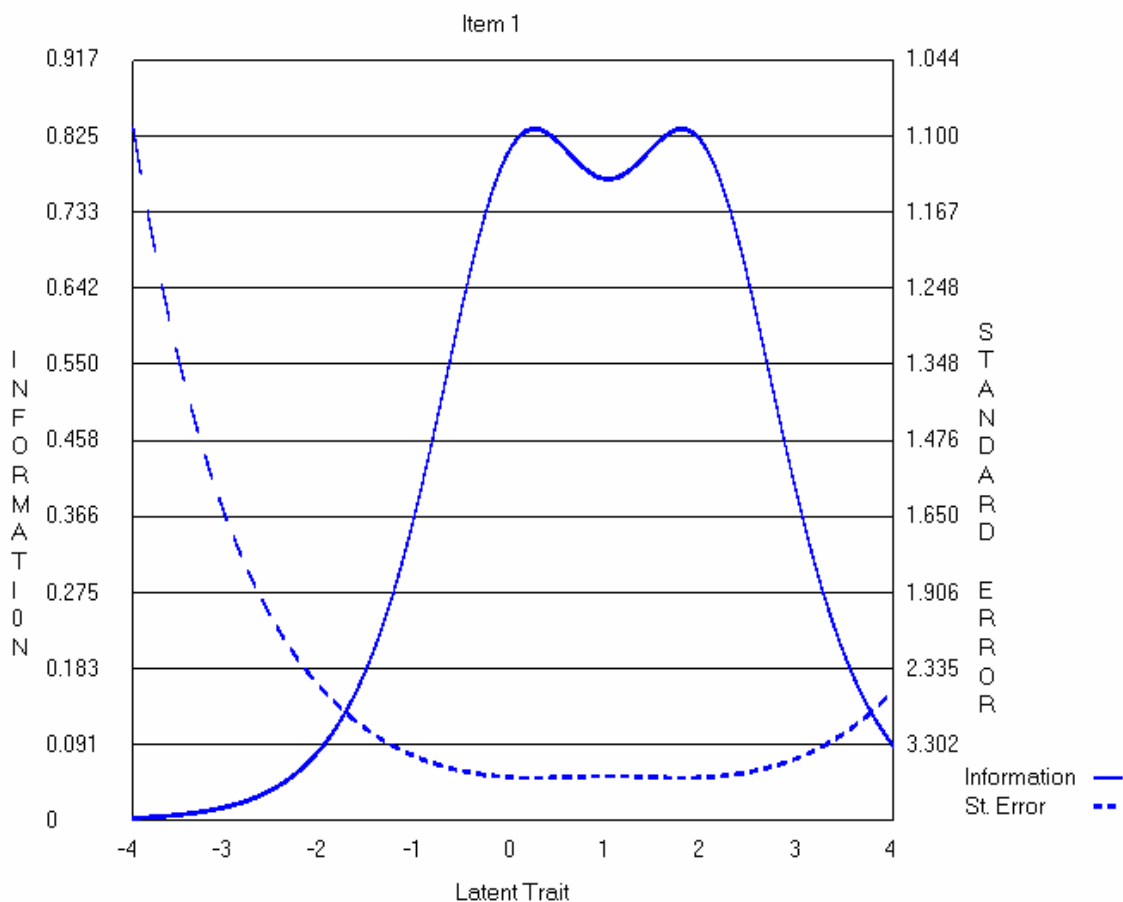
View Information Functions

The View Information Functions Window presents the item information function for any item of the test constructed in the Main Window, as well as the test information function for the test constructed in the Main Window. The user can display the information function for any item or test by left-clicking on the appropriate row of the drop-down list at the top of the View Information Functions Window. The information function for the related item or test will appear in the image at the bottom of the View Information Functions Window. An example of the resulting image for a polytomous item following the GPCM having a = 1, b1 = 0, b2 = 2 is shown below.

Item 1

Latent Trait	
0.917	
0.825	
0.733	
0.642	
0.550	
0.458	
0.366	
0.275	
0.183	
0.091	
0	

-4 -3 -2 -1 0 1 2 3 4

Latent Trait

The resulting image can be copied and pasted into other applications by clicking on the "Copy" button at the right side of the View Information Functions Window.

Compare Information and SE

Because the standard error of the theta estimate is the square root of the inverse of the information function, it is often of interest to examine the information function and the standard error simultaneously. The Compare Information and SE Window allows the user to do this by selecting the appropriate item or test from the drop-down list at the top of the Compare Information and SE Window. The resulting image displays the information function and the conditional standard error. The values of the information function are presented on the left side of the figure, and the values of the standard error are presented on the right side of the figure. An example of the resulting image for a polytomous item following the GPCM having a = 1, b1 = 0, b2 = 2 is shown below.

Item 1

INFORMATION	value		STANDARD ERROR	value

The resulting image can be copied and pasted into other applications by clicking on the "Copy" button at the right side of the Compare Information and SE Window.

CHAPTER 4
Simulate Data Window

The Simulate Data Window permits the user to simulate responses to the test specified in the Main Window. This section provides instructions for using the Simulate Data Window

Simulating Data

To simulate data using the Simulate Data Window of IRT-Lab, follow the steps outlined below.

Step 1: Construct a test using the Main Window.

Step 2: In the Simulate Data Window, specify the number of cases for which you wish to simulate responses to the items of the test. The number of cases specified must be equal to or greater than one. Note that if you are simulating data for multiple groups, you must specify the number of cases associated with each group.

Step 3: In the Simulate Data Window, specify the mean and standard deviation of the normal distribution of theta. By default, the mean is set to zero and the standard deviation to one. If you are simulating data for multiple groups, you may assign different means and standard deviations for each group.

Step 4: In the Simulate Data Window, specify the number of trials to run in the simulation. By default, the number of trials is set to one. However, you can specify any number of trials you wish.

Step 5: In the Simulate Data Window, specify the name of the text file to which the simulated responses will be saved, and the location of the text file.

Step 6: In the Simulate Data Window, specify whether the observations of the saved file are delimited by a space, a comma, or a tab. By default, the saved simulated responses are delimited by a space.

Step 7: Click the "Simulate" button. The resulting simulated responses are located in the text file(s) specified in the lower portion of the Simulate Data Window.

If more than one trial is selected, then the file name used to save the resulting text files are set to "FilenameTrial#.txt". For example, if the file name specified by the user is

"Simulation", then the resulting file names for trials 1, 2, 3, etc., are
"SimulationTria1.txt", "SimulationTria2.txt", "SimulationTrial3.txt", etc.

Example 1: Simulating Responses for a Single Group

Consider an example of a ten-item test for which a researcher wishes to simulate
responses for a random sample of 20 individuals drawn from a single group. The ten
items of the test have the following properties (as obtained from the saved text log file
using IRT-Lab):

```
Item   Grp    Mod    c      a      b1      b2      b3      b4      b5      b6
1      1      3PL    .2     1      -1.5
2      1      3PL    .2     1      -1.5
3      1      3PL    .2     1      -1.5
4      1      3PL    .2     1      0
5      1      3PL    .2     1      0
6      1      3PL    .2     1      0
7      1      3PL    .2     1      0
8      1      3PL    .2     1      1.5
9      1      3PL    .2     1      1.5
10     1      3PL    .2     1      1.5
```

Thus, the test consists of ten dichotomously scored items that follow the 3PL. Using this
test, responses from a random sample of 20 individuals (for which θ follows a normal
distribution with mean 0 and standard deviation 1) are simulated. The resulting text file
(space-delimited) containing the simulated responses is given by:

```
1 1 1 1 1 1 1 0 0 0
1 1 1 1 0 1 1 0 0 1
1 1 1 1 0 1 1 1 0 1
1 1 1 1 1 0 1 1 0 0
1 1 1 1 0 0 1 1 1 0
1 1 1 0 0 1 1 0 1 0
1 0 1 1 1 1 0 0 0 0
1 1 1 1 1 0 0 0 0 1
1 1 0 0 1 0 1 0 0 0
1 1 1 1 1 1 1 0 1 0
1 1 1 1 0 1 1 0 0 0
1 1 1 0 1 1 0 0 0 0
1 1 1 0 0 1 0 0 0 1
1 1 1 1 1 1 1 0 1 0
1 1 1 1 1 1 1 1 0 0
1 1 1 1 1 1 1 0 0 0
1 1 1 1 1 1 1 0 0 0
1 1 1 1 1 1 1 0 0 0
1 1 1 1 0 1 1 0 0 1
0 1 1 0 1 0 0 1 1 1
```

Each row corresponds to the responses of a single individual, and each column
corresponds to the responses of the individuals to a single item. Because the Main Page
consisted of parameters for the items of a test for only a single group, no group identifier
has been included in the outputted file containing the simulated responses.

As a pedagogical tool, notice that the first three items were relatively easy (b1 = -1.5), the middle four items were of medium difficulty (b1 = 0) and the last three items were relatively difficulty (b1 = 1.5). Examination of the responses you shows that the responses are consistent with these levels of difficulty; the responses to the first three items were largely "1", the responses to the middle four items were a combination of "0" and "1", and the responses to the last three items were largely "0".

Example 2: Simulating Responses for Multiple Groups

Consider an example of a five-item test for which item parameters are specified for two different groups of examinees. The five items of the test for each of the two groups have the following properties (as obtained from the saved text log file using IRT-Lab):

```
Item    Grp    Mod      c      a      b1      b2      b3      b4      b5      b6
1       1      Rasch    0      1      -1
1       2      Rasch    0      1      1
2       1      Rasch    0      1      -1
2       2      Rasch    0      1      1
3       1      Rasch    0      1      -1
3       2      Rasch    0      1      1
4       1      Rasch    0      1      -1
4       2      Rasch    0      1      1
5       1      Rasch    0      1      -1
5       2      Rasch    0      1      1
```

Notice that each of the five items has a difficulty parameter (b1) equal to −1 for Group 1 and 1 for Group 2. Thus, the items of the test are more difficulty for Group 2 than Group 1. Using this test, responses from a random sample of 10 individuals from each group were simulated. The resulting simulated responses (obtained from the saved space-delimited text file) are as follows.

```
1 1 1 1 1 1
1 0 0 1 1 1
1 1 1 1 1 1
1 1 1 1 1 1
1 0 1 0 1 0
1 0 0 0 0 0
1 0 0 1 0 0
1 0 1 1 1 1
1 1 1 1 1 1
1 0 1 0 1 1
2 0 0 0 0 0
2 1 0 1 1 0
2 0 1 1 1 1
2 0 0 0 0 0
2 0 1 0 0 0
2 1 0 1 1 0
2 0 0 0 0 0
2 0 0 1 1 0
2 1 0 1 0 0
2 1 1 1 0 0
```

Notice that because there are two groups, the first column corresponds to a group identifier. In addition, because the value of the b1-parameter was lower for Group 1 than for Group 2, you will find that there are more correct responses (1) for Group 1 than for Group 2.

Simulation Procedures

The simulation of responses is conducted via the following steps.

1. For a given individual in a simulated sample, a latent trait value (θ) is randomly assigned according to a normal distribution having mean and standard deviation specified in the Simulate Data Window.

2. For the individual in question, a simulated response for a given item is obtained by:

> (i) Obtaining the probability of observing response options $j = 0, 1, \ldots, m$, for the item in question given the individual's value of θ. The relevant probabilities are obtained using the equations of the ICCs as outlined in Chapter 5. Note that for dichotomous models (Rasch, 1PL, 2PL, and 3PL) the response options consist of $j = 0, 1$ only, and for polytomous models (GPCM and GRM) the response options consist of $j = 0, 1, \ldots, m$. The probability of response option j will be denoted here by P_j.

> (ii) Using the probabilities of each response option (P_j) obtained in (i) above, the cumulative probability of each response option (γ_j) is obtained, such that $\gamma_0 = P_0$, $\gamma_1 = P_0 + P_1, \ldots, \gamma_m = P_0 + P_1 + \ldots + P_m$.

> (iii) A random uniform variate on the interval (0, 1) is obtained. This uniform variate is denoted here by u.

> (iv) Using the cumulative probabilities and the random uniform variate obtained in (ii) and (iii) above, the lowest value of j for which $\gamma_j \geq u$ is found. The obtained value of j is assigned to the simulated response for the item.

3. Step 2 is repeated for each item of the test. The responses to the test items for a given individual are reported in a single line of the text file containing the simulated responses. If there is more than one group, the first value in each row specifies the group number. If only one group exists, then a group identifier is not present.

4. Steps 1, 2 and 3 are repeated for each individual in each of the simulated samples (one sample is drawn for each group for which the test is defined).

5. Steps 1, 2, 3, and 4 are repeated for each trial in the simulation.

CHAPTER 5
Equations of IRT-Lab

This section outlines the equations used to compute the ICCs and information functions employed in IRT-Lab.

Item Characteristic Curves

Consider a response variable Y, with response options represented by $y = 0, 1, \ldots, m$. In the case of a dichotomous item, the incorrect response is represented by $y = 0$ and the correct response by $y = 1$. For the models used in IRT-Lab, the model parameters are defined as follows:

θ The latent construct being measured by the item
b_y The location parameter associated with response y
a The discrimination parameter
c The guessing parameter
D The normalizing coefficient of 1.702

The Rasch model specifies the probability of correct response by

$$P(Y = 1 \mid \theta) = \frac{\exp[\theta - b_1]}{1 + \exp[\theta - b_1]}.$$

The one-parameter logistic model (1PL) specifies the probability of correct response by

$$P(Y = 1 \mid \theta) = \frac{\exp[D(\theta - b_1)]}{1 + \exp[D(\theta - b_1)]}.$$

The two-parameter logistic model (2PL) specifies the probability of correct response by

$$P(Y = 1 \mid \theta) = \frac{\exp[Da(\theta - b_1)]}{1 + \exp[Da(\theta - b_1)]}.$$

The three-parameter logistic model (3PL) specifies the probability of correct response by

$$P(Y = 1 \mid \theta) = c + (1 - c) \frac{\exp[Da(\theta - b_1)]}{1 + \exp[Da(\theta - b_1)]}.$$

The generalized partial credit model (GPCM) specifies the probability of observing response option y by

$$P(Y = y) = \frac{\exp\left[\sum_{j=0}^{y} Da(\theta - b_j)\right]}{\sum_{j=0}^{m}\left[\exp \sum_{k=0}^{j} Da(\theta - b_k)\right]}.$$

The partial credit model (PCM) is obtained by setting the a-parameter equal to

$$a = \frac{1}{D} = 0.588.$$

The graded response model (GRM) specifies the probability of observing response option y by

$$P(Y = y \mid \theta) = P(Y \geq y) - P(Y \geq y + 1)$$

where

$$P(Y \geq y \mid \theta) = \frac{\exp[Da(\theta - b_y)]}{1 + \exp[Da(\theta - b_y)]} \quad \text{for } y > 0$$

and

$$P(Y \geq y \mid \theta) = 1 \quad \text{for } y = 0.$$